



National Security Agency/Central Support Service



INFORMATION ASSURANCE DIRECTORATE

CGS Network Boundary Protection Capability

Version 1.1.1

Network Boundary Protection is the Capability to protect and control access to Enterprise resources across a security boundary. A security boundary exists when there is a separation of entities (systems, networks, enclaves, or Enterprises), which are governed by differing security policies or operate in a different threat environment. Network Boundary Protection is carried out by placing information assurance (IA) mechanisms between the internal system and the systems external to the security boundary.



CGS Network Boundary Protection Capability



Version 1.1.1

Table of Contents

1	Revisions	2
2	Capability Definition	3
3	Capability Gold Standard Guidance.....	4
4	Environment Pre-Conditions	8
5	Capability Post-Conditions.....	8
6	Organizational Implementation Considerations	8
7	Capability Interrelationships.....	11
7.1	Required Interrelationships	11
7.2	Core Interrelationships	13
7.3	Supporting Interrelationships.....	13
8	Security Controls	14
9	Directives, Policies, and Standards	17
10	Cost Considerations	20
11	Guidance Statements	21



CGS Network Boundary Protection Capability



Version 1.1.1

1 Revisions

Name	Date	Reason	Version
CGS Team	30 June 2011	Initial release	1.1
CGS Team	30 July 2012	Inclusion of new IAD document template & Synopsis	1.1.1



CGS Network Boundary Protection Capability



Version 1.1.1

2 Capability Definition

The Capability definition provides an understanding of the importance of the Capability to the Enterprise. It provides a high-level overview of the Capability based on definitions derived from Committee on National Security Systems (CNSS) Instruction No. 4009.

Network Boundary Protection is the Capability to protect and control access to Enterprise resources across a security boundary. A security boundary exists when there is a separation of entities (systems, networks, enclaves, or Enterprises), which are governed by differing security policies or operate in a different threat environment. Network Boundary Protection is carried out by placing information assurance (IA) mechanisms between the internal system and the systems external to the security boundary. Examples of such IA mechanisms include, but are not limited to, Cross-Domain Solutions (CDSs), Controlled Interfaces, Demilitarized Zones (DMZs), Virtual Private Networks (VPNs), and encryption devices at the boundary, or interface. Because of the differing nature of boundary protection devices, a brief explanation of the example technologies is provided here for completeness.

The purpose of a CDS is to provide a manual or automated means for transferring data between two or more differing security domains. The CDS is responsible for ensuring that unauthorized information cannot leak from a domain with information controlled at a higher level (e.g., classified information) to a domain that is controlled at a lower level (e.g., unclassified information). The CDS is also responsible for protecting the network from malicious content that may be passed from the less controlled network. CDSs enable the transfer of information among security domains, which is normally prohibited by automated policies, but is required for successful completion of a mission.

The purpose of a Controlled Interface is to control access and information flow into and out of the domain. A Controlled Interface is used when the security policy between interconnected domains is fairly similar. In this instance, the networks are at the same classification level and the risk of contamination or attack by another domain is considered to be sufficiently low.

The purpose of a DMZ is to provide an additional layer of security to an Organization's network when some of its services must be exposed to a larger community. A DMZ can be a physical or logical subnetwork. With a DMZ, the Organization can provide an external face and external services, while controlling interactions with the internal network



CGS Network Boundary Protection Capability



Version 1.1.1

and ensuring that an external attacker is restricted to the equipment in the DMZ, rather than having access to any part of the internal network.

The purpose of a secure VPN (not a traffic engineering VPN) is to provide a connection into the network from a remote point, either for a user to gain access or to establish a connection between two networks. To perform this function, the VPN establishes a relationship between the endpoints (through authentication mechanisms) and then encrypts traffic between those endpoints such that traffic is protected from the underlying network.

3 Capability Gold Standard Guidance

The Capability Gold Standard Guidance evaluates the Enterprise needs and overlays the expected Gold Standard behavior. The guidance goes beyond the concept of “good enough” when describing the Gold Standard recommendations, considers industry best practices, and describes a level of security that not only meets current standards but also exceeds them across the Enterprise.

The first step in securing a boundary is to understand where the boundary is, the security environment and policy for each side of the boundary, and the mission the boundary is supporting. This understanding is provided by the Network Boundary and Interfaces and Understand Mission Flow Capabilities. It is also key to understanding and defining the trust relationship between the connecting enclaves to determine what level of protection is needed at each boundary. Based on this knowledge, each network used by the Enterprise shall have its boundaries secured to enforce confidentiality, integrity, availability, authentication, non-repudiation, and access control rules, thereby providing essential protection to authorized users, missions, and data to meet the specific needs of the boundary. Although a boundary is defined by internal versus external, there may be instances where a boundary is considered internal to the overall Enterprise, as defined by the agency/Organization. For example, there could be an internal boundary between enclaves within an agency in the Intelligence Community (IC), Department of Defense (DoD), or civil agency. An external connection would occur when the systems on the outside of the boundary are out of the control of the internal network. For example, a connection between a DoD agency and an IC agency would be considered an external connection.

All network boundaries shall (based on connection type, mission needs, risk, etc.) incorporate approved firewalls, intrusion prevention/detection solutions,



CGS Network Boundary Protection Capability



Version 1.1.1

spam/malware/content filtering, DMZs for vulnerable border resources, and/or CDSs where the network connects to another network of a different security policy or environment. Approved network boundary devices are defined as those devices that have been certified and accredited for use in that environment, based on the security requirements enforced and the threat environment. The Enterprise shall employ vendor diversity or defense-in-depth strategies for devices, where appropriate. If remote connections are authorized on the network, they shall be verified individually each time a connection is attempted because they are prime targets for attack. Physical access to network resources shall also be secured by whatever means necessary (as deemed by the accrediting authority), including investigating users prior to granting access, authenticating users via secure credentials upon entry, employing guards, using security systems. All local resources shall be verified, local connections authorized, and wireless access points (WAPs) secured. System/network availability and throughput requirements shall be taken into account to ensure that Network Boundary Protection services meet these requirements and policies. See the Organizational Implementation Considerations section for more information.

The boundary devices shall provide sufficient levels of protection at all layers (i.e., physical through application) to ensure the added risks are sufficiently mitigated. Processes shall be in place to monitor both the boundary defenses and all traffic flowing through it. When a network boundary device detects an event or incident, the device shall forward the incident information to the Incident Response Capability within an Organization-defined amount of time as determined by mission and risk needs. Because of the broad and differing nature of the mechanisms that make up the Network Boundary Protection Capability, the following paragraphs provide additional information for the example mechanisms provided in the definition section.

Cross Domain Solutions

Similar to Controlled Interfaces, CDSs are used to perform bidirectional filtering functions that prevent data leakage and stop the passing of malicious software. CDSs can filter traffic at varying layers of the Open System Interconnection (OSI) stack. Because CDSs are specialized devices, it is important to understand the specific boundary's data transfer needs and choose a solution that meets those specific needs. Because CDSs are placed at the boundary of enclaves with differing security domains, these devices shall be high-assurance devices (i.e., devices that have been carefully designed and implemented for high-risk situations) and undergo specialized certification testing. These test results are then used, in conjunction with knowledge of the data type and connecting enclaves, to make a risk decision and accredit the solution. The Organizations shall work with their



CGS Network Boundary Protection Capability



Version 1.1.1

appropriate Cross- Domain Support Office (CDSO) or their governing Organization responsible for CDS decisions to ensure they select the appropriate solution. See the Organizational Implementation Considerations section for further information regarding product implementation decisions. When dictated by mission need, the Enterprise shall employ CDSs that enable the following functions:

- Provision of centralized and remote management (coordinated with Communication Protection)
- Provision of cross-domain chat
- Sharing of high-risk files (e.g., Microsoft Office files, PDFs, executables)
- Discovery searches of low-side servers by high-side users (including file retrieval)
- Remote management of low-side infrastructure devices from the high side
- Enhanced cross-domain flow of situation awareness information

Controlled Interface

Controlled Interfaces typically include a secure gateway, firewall, filtering routers, application intelligence or proxies, or deep packet, stateful inspection filtering functions. These devices are capable of bidirectionally screening or filtering information and preventing data leakage as well as blocking unallowed traffic. Malware detection software is used to detect known viruses and Trojan horses, among others. Intrusion detection systems are used to detect attacks. In addition, Controlled Interfaces may provide secure transport from the protected network to the interface point and reliable, redundant network access service to protect against denial of service or other such attacks. This functionality may be performed at a single interface point or may be distributed across multiple Controlled Interfaces.

Demilitarized Zones

In a network, the hosts that are the most vulnerable to attack are those that provide services to users outside of the local network such as email, web, and domain name service (DNS) servers. Because of the increased potential of these servers being compromised, they are placed into their own subnetwork to protect the rest of the network in case an intruder is successful. In addition, these hosts are security hardened (see the System Protection Capability) to limit the susceptibility to attacks from external networks. Hosts in the DMZ have limited connectivity to specific hosts in the internal network, though communication with other nodes in the DMZ and with the external network is allowed. This allows servers in the DMZ to provide services to both the internal and external networks, while an intervening firewall controls the traffic between the DMZ servers and the external network clients. DMZs shall be placed as close to the network boundary as possible and implemented with specialized IA controls and devices that are



CGS Network Boundary Protection Capability



Version 1.1.1

specifically designed for the protocols or services used at the boundary connection. For example, an email proxy can be placed in front of an email server, and firewalls ensure that backend services are blocked (i.e., a user needing access to specific data shall be provided access to the web service for that data, not the database itself).

Virtual Private Networks

Many networks shall be able to support remote access, and in some cases, this remote access is to an internal system. Remote access shall be provided by a secure service, such as a VPN, that provides two-way authentication between the remote user and local system and, at a minimum, encrypts the data when outside the network boundary. VPN usage can also include site-to-site implementations. In this case, it is important to ensure encrypted traffic is passed only to allowed access points and the appropriate policies are defined for the encrypted traffic (e.g., it is dropped at the boundary). It is important to understand the type of access and functionality (real-time application execution versus user email access) that needs to be allowed with each remote connection. Understanding this ensures that the risk of allowing remote access is understood, and the data and network are protected accordingly when the risk is acceptable. For example, it may be important to ensure that data cannot reside on the user's local system.

All boundary devices shall be capable (in accordance with the defined policy for that boundary) of:

- Performing bidirectional information services or data flows
- Filtering, at the Internet Protocol (IP) (routes) and application (deep packet) layer, as well as packet inspection
- Scanning for malicious content
- Authenticating traffic sources and destinations and enforcing the access policy
- Auditing, monitoring, and responding or providing information to the appropriate Community Gold Standard Capability to respond
- Meeting the Enterprise's throughput and availability requirements, providing mechanisms such as redundancy and failover, where appropriate
- Enforcing policies set by the Port Security Capability

There shall be trustworthy mechanisms to administer, manage, and monitor Network Boundary Protection solutions. The devices shall be able to securely connect to the operations center using secure connections as defined within the Communications Security Capability and be configured securely as defined within the System Protection Capability. All devices shall be on the evaluated and approved products list for the Organization and the specific environment where they are being implemented.



CGS Network Boundary Protection Capability



Version 1.1.1

4 Environment Pre-Conditions

The environment pre-conditions provide insight into environmental, user, and technological aspects needed for Capability implementation. These pre-conditions are services or other Capabilities that must be in place within the Enterprise for the Capability to function.

1. The environment knows what ports, protocols, and allowed IP ranges and services are authorized to be forwarded.
2. The Enterprise will have identified each of its internal and external boundaries.
3. Attackers may be in the Enterprise, and insider threat exists.
4. WAPs, and other network boundary devices, can ensure they are providing service only to authorized devices.
5. The Enterprise's infrastructure is physically secure.
6. The environment provides boundary components that are configured securely.
7. Monitoring, reporting, auditing, and preventing and detecting intrusion are handled in the environment.
8. Redundant/failover systems are configured identically.
9. Peer boundaries are configured with the same protections.

5 Capability Post-Conditions

The Capability post-conditions define what the Capability will provide. They define functions that the Capability will perform or constraints that the Capability will operate under when performing its function.

1. The network boundary policies will be configured according to security best practices (i.e., deny all, explicit allow).
2. Peer boundaries will be configured with the same security policies.
3. The Capability provides traffic scanning for malicious activity.
4. The Capability provides content filtering.
5. The Capability maintains necessary system uptime and throughput requirements.
6. The Enterprise will use Network Boundary Protection devices, such as CDSs, Controlled Interfaces, DMZs, and VPNs, at all security boundaries.

6 Organizational Implementation Considerations

Organizational implementation considerations provide insight into what the Organization needs to establish, ensure, and have in place for the specified Capability to be effective. It provides guidance specific to the actions, people, processes, and departments that an Organization will need to execute or establish to implement the guidance described in Section 3 (Capability Gold Standard Guidance).



CGS Network Boundary Protection Capability



Version 1.1.1

When Network Boundary Protection is implemented correctly, the department or agency will possess a Capability to prevent unauthorized access to its networks and provide secure usage to its authorized users. Potentially vulnerable border resources, such as web servers, will be secured inside of DMZs. CDSs will be used to connect with networks that carry a different security level; Controlled Interfaces and other boundary devices will be used when connecting enclaves, as appropriate. VPNs will be used when remote access is required.

Physical access to the network boundary devices will be restricted to only those users who have been investigated (if applicable) and authorized. Physical protections will include authentication via secure credentials, guards, checkpoints, security systems, etc.

In general, Network Boundary Protection is made up of a suite of IA devices as opposed to a single device. These devices are responsible for controlling the inbound and outbound flow of information in accordance with the policies they implement. The boundary devices will be configured in accordance with security guidance and best practices and will also maintain availability requirements in accordance with the Enterprise needs and policy, such that the Organization will define a fail open or fail closed policy or require rollover or redundancy on a case-by-case decision. These requirements will also lead to the determination of the actual boundary protection mechanism the Enterprise will implement. When selecting specific mechanisms, the Organization will be responsible for ensuring that the mechanisms can communicate via standard protocols, where possible. In addition, when selecting devices, the Organization will be cognizant of vendor selection and ensure it meets acquisition requirements.

When connecting one environment to another environment at the same classification level, the Organization will not make the assumption that the risk and threat are the same. The Organization will need to decide what the risk is to its environment and make implementation decisions based on those risk decisions. Regardless, when implementing any boundary device, the Organization will need to ensure the products and connections obtain appropriate approval prior to being introduced into the environment. Decisions regarding the type of boundary device will be based on data, need to know, likelihood of attack, and partner type (connecting environments), at a minimum.

Cross-Domain Solutions

CDSs are high-assurance devices that go through specialized certification testing and a specific accreditation process. Because CDSs are used to transfer data across security



CGS Network Boundary Protection Capability



Version 1.1.1

(e.g., classification) domains, the Organization will work with its designated CDSO or governing office to ensure the connection requirements are clearly understood and the appropriate approval processes are followed. For example, the Organization will be able to articulate the data type it needs to transfer across security domains and the networks to which it is going to connect. Based on the defined requirements, the CDS governing office will assist the Organization in determining whether its specific data transfer needs can be met by an existing Enterprise service or CDS on the Unified Cross-Domain Management Office (UCDMO) Baseline. If the Organization's needs cannot be met, the CDS office will help the Organization to determine its next steps: whether to modify an existing product or develop a new product. The key to implementing a CDS is to understand the risk the solution poses to the network and to ensure the implementation is approved by the appropriate Designated Approval Authority (DAA).

Controlled Interface

Connections between networks or other enclaves and the core will be few in number (based on mission requirements, as defined in the Understand Mission Flow Capability), tightly controlled, and provide the required access with the necessary security. Features that are required, based on mission need and risk, will include a firewall/filtering (i.e., filtering routers, application intelligence, or a proxy) system that supports the element's highest priority applications, an intrusion detection system, secure transport from the component network to the Controlled Interface, redundant access service, and supply chain risk management provisions.

Demilitarized Zones

There are many different ways to design a network with a DMZ. Two of the most basic methods are using a single firewall, also known as the three-legged model, and using dual firewalls. These architectures can be expanded to create very complex architectures, depending on the network requirements. A single firewall with at least three network interfaces can be used to create a DMZ. One interface connects to the external network, a second interface connects to the DMZ servers, and the third interface connects to the internal network. All traffic passes through and is controlled by the firewall. This architecture works well for small systems with low traffic, but it does not scale well because it creates a single point of failure (if the firewall fails, the connection fails; if the firewall is defeated, the security provided by the DMZ is negated). In addition, this architecture may be acceptable for an internal boundary but not for an external boundary, depending on risk.



CGS Network Boundary Protection Capability



Version 1.1.1

A more scalable and secure architecture uses two firewalls to create a DMZ. The first firewall sits between the DMZ and the external network and filters all traffic coming from and going to the external network. The second firewall sits between the DMZ resources and the internal network. This provides a second level of protection for the internal network from external attacks and also provides a line of defense for attacks originating on the DMZ's subnet (e.g., if the DMZ has been successfully compromised). The use of two firewalls also provides two levels of defense against attacks that could originate on the internal network (e.g., an insider is prevented from using this network as a base from which to attack other resources). To maximize defenses, the firewalls are diversified (use of multiple vendors).

Virtual Private Networks

For those networks that allow remote connections, those connections all need to be verified individually every time a connection is attempted. VPNs can be connections between an individual user and an internal system or site-to-site connections. The Organizations will define the appropriate policies for the connection depending on the type of connection and access being provided. For a connection to be authorized, the user will need to be authenticated by supplying valid credentials. Depending on policy, as determined by the Organization, the connection source will need to be verified, in which case an authorized user trying to connect on a borrowed laptop may not be allowed. VPNs may be implemented using a VPN concentrator. In which case, the concentrator will be housed within a DMZ.

7 Capability Interrelationships

Capability interrelationships identify other Capabilities within the Community Gold Standard framework that the Capability in this document relies on to operate. Although there are many relationships between the Capabilities, the focus is on the primary relationships in which the Capabilities directly communicate with or influence one another.

7.1 Required Interrelationships

The following Capability interrelationships include the other Capabilities within the Community Gold Standard framework that are necessary for the Capability in this document to operate.

- Network Mapping—The Network Boundary Protection Capability relies on the Network Mapping Capability to provide information about the network that is used in determining protection requirements.



CGS Network Boundary Protection Capability



Version 1.1.1

- Network Boundary and Interfaces—The Network Boundary Protection Capability relies on the Network Boundary and Interfaces Capability to identify entry and exit points for a network.
- Understand Mission Flows—The Network Boundary Protection Capability relies on the Understand Mission Flows Capability for information about the protection requirements for Enterprise resources.
- Understand Data Flows—The Network Boundary Protection Capability relies on the Understand Data Flows Capability for information about the protection requirements for Enterprise resources.
- System Protection—The Network Boundary Protection Capability relies on the System Protection Capability to provide protection mechanisms for network boundary protection devices.
- Communication Protection—The Network Boundary Protection Capability relies on the Communication Protection Capability to control and secure data flows across network boundaries.
- Configuration Management—The Network Boundary Protection Capability relies on the Configuration Management Capability to ensure that network boundary protection devices are compliant with configurations as outlined in the Configuration Management Plan.
- Identity Management—The Network Boundary Protection Capability relies on the Identity Management Capability to ensure that only resources with known and approved identities are entering and exiting the network. All network boundary devices must have identity management controls enforced so that only uniquely or appropriately identified personnel/devices are able to administer the network boundary device.
- Access Management—The Network Boundary Protection Capability relies on the Access Management Capability to ensure that only authorized resources and traffic are entering and exiting the network. In addition, the Network Boundary Protection Capability relies on the Access Management Capability to ensure that all network boundary devices have access controls enforced so that only authorized personnel/devices are able to administer the network boundary device.
- Metadata Management—The Network Boundary Protection Capability relies on Metadata Management to ensure that correct security labels are bound to data to be passed across the boundary so that the inspection and protection mechanisms make the correct decision regarding release.
- Architecture Reviews—The Network Boundary Protection relies on the Architecture Reviews Capability to assess the security controls of a system to ensure that IA



CGS Network Boundary Protection Capability



Version 1.1.1

concepts (e.g., confidentiality, integrity, availability, authentication, and non-repudiation) are present in Enterprise architecture requirements.

- Network Intrusion Detection—The Network Boundary Protection Capability relies on the Network Intrusion Detection Capability to manage network intrusion detection devices that operate as a part of the protection mechanisms at network boundaries.
- Network Intrusion Prevention—The Network Boundary Protection Capability relies on the Network Intrusion Prevention Capability to manage network intrusion prevention devices that operate as a part of the protection mechanisms at network boundaries.

7.2 Core Interrelationships

The following Capability interrelationships include the Capabilities within the Community Gold Standard framework that relate to every Capability.

- Portfolio Management—The Network Boundary Protection Capability relies on the Portfolio Management Capability to determine current and future investment needs and prioritize investments based on those needs.
- IA Policies, Procedures, and Standards—The Network Boundary Protection Capability relies on the IA Policies, Procedures, and Standards Capability to provide information about applicable federal laws, Executive Orders, regulations, directives, policies, procedures, and standards.
- IA Awareness—The Network Boundary Protection Capability relies on the IA Awareness Capability for an awareness program to inform personnel of their responsibilities related to IA.
- IA Training—The Network Boundary Protection Capability relies on the IA Training Capability to provide training programs related to IA activities in accordance with agency policies.
- Organizations and Authorities—The Network Boundary Protection Capability relies on the Organizations and Authorities Capability to establish the relevant roles and responsibilities.

7.3 Supporting Interrelationships

The following Capability interrelationships include the other Capabilities within the Community Gold Standard framework that are not necessary for the Capability to operate, although they support the operation of the Capability in this document.

- Risk Analysis—The Network Boundary Protection Capability establishes protection mechanisms that are part of an accredited system and documented as such



CGS Network Boundary Protection Capability



Version 1.1.1

through a certification and accreditation process conducted by the Risk Analysis Capability.

- Risk Mitigation—The Network Boundary Protection Capability implements individual countermeasures that may be selected by the Risk Mitigation Capability.

8 Security Controls

This section provides a mapping of the Capability to the appropriate controls. The controls and their enhancements are granularly mapped according to their applicability. In some instances, a control may map to multiple Capabilities.

Control Number/Title	Related Text
NIST SP 800-53 Revision 3, <i>Recommended Security Controls for Federal Information Systems and Organizations</i>	
AC-4 INFORMATION FLOW ENFORCEMENT	Enhancement/s: (10) The information system provides the capability for a privileged administrator to enable/disable [Assignment: organization-defined security policy filters]. (11) The information system provides the capability for a privileged administrator to configure [Assignment: organization-defined security policy filters] to support different security policies. (12) The information system, when transferring information between different security domains, identifies information flows by data type specification and usage. (13) The information system, when transferring information between different security domains, decomposes information into policy-relevant subcomponents for submission to policy enforcement mechanisms. (14) The information system, when transferring information between different security domains, implements policy filters that constrain data structure and content to [Assignment: organization-defined information security policy requirements]. (15) The information system, when transferring information between different security domains, detects unsanctioned information and prohibits the transfer of such information in accordance with the security policy.
CA-3 INFORMATION	Control: The organization:



CGS Network Boundary Protection Capability



Version 1.1.1

SYSTEM CONNECTIONS	<p>a. Authorizes connections from the information system to other information systems outside of the authorization boundary through the use of Interconnection Security Agreements;</p> <p>b. Documents, for each connection, the interface characteristics, security requirements, and the nature of the information communicated; and</p> <p>c. Monitors the information system connections on an ongoing basis verifying enforcement of security requirements.</p> <p>Enhancement/s:</p> <p>(1) The organization prohibits the direct connection of an unclassified, national security system to an external network.</p> <p>(2) The organization prohibits the direct connection of a classified, national security system to an external network.</p>
SC-7 BOUNDARY PROTECTION	<p>Control: The information system:</p> <p>a. Monitors and controls communications at the external boundary of the system and at key internal boundaries within the system;</p> <p>b. Connects to external networks or information systems only through managed interfaces consisting of boundary protection devices arranged in accordance with an organizational architecture.</p> <p>Enhancement/s:</p> <p>(1) The organization physically allocates publicly accessible information system components to separate subnetworks with separate physical network interfaces.</p> <p>(2) The information system prevents public access into the organization's internal networks except as appropriately mediated by managed interfaces employing boundary protection devices.</p> <p>(3) The organization limits the number of access points to the information system to allow for more comprehensive monitoring of inbound and outbound communications and network traffic.</p> <p>(4) The organization:</p> <p>(a) Implements a managed interface for each external telecommunication service;</p> <p>(b) Establishes a traffic flow policy for each managed interface;</p> <p>(c) Employs security controls as needed to protect the confidentiality and integrity of the information being transmitted;</p>



CGS Network Boundary Protection Capability



Version 1.1.1

(d) Documents each exception to the traffic flow policy with a supporting mission/business need and duration of that need;

(e) Reviews exceptions to the traffic flow policy [Assignment: organization-defined frequency]; and

(f) Removes traffic flow policy exceptions that are no longer supported by an explicit mission/business need.

(5) The information system at managed interfaces, denies network traffic by default and allows network traffic by exception (i.e., deny all, permit by exception).

(6) The organization prevents the unauthorized release of information outside of the information system boundary or any unauthorized communication through the information system boundary when there is an operational failure of the boundary protection mechanisms.

(7) The information system prevents remote devices that have established a non-remote connection with the system from communicating outside of that communications path with resources in external networks.

(8) The information system routes [Assignment: organization-defined internal communications traffic] to [Assignment: organization-defined external networks] through authenticated proxy servers within the managed interfaces of boundary protection devices.

(9) The information system, at managed interfaces, denies network traffic and audits internal users (or malicious code) posing a threat to external information systems.

(10) The organization prevents the unauthorized exfiltration of information across managed interfaces.

(11) The information system checks incoming communications to ensure that the communications are coming from an authorized source and routed to an authorized destination.

(12) The information system implements host-based boundary protection mechanisms for servers, workstations, and mobile devices.

(13) The organization isolates [Assignment: organization defined key information security tools, mechanisms, and support components] from other internal information system components via physically separate subnets with managed interfaces to other portions of the system.



CGS Network Boundary Protection Capability



Version 1.1.1

	<p>(14) The organization protects against unauthorized physical connections across the boundary protections implemented at [Assignment: organization-defined list of managed interfaces].</p> <p>(15) The information system routes all networked, privileged access through a dedicated, managed interface for purposes of access control and auditing.</p> <p>(16) The information system prevents discovery of specific system components (or devices) composing a managed interface.</p> <p>(17) The organization employs automated mechanisms to enforce strict adherence to protocol format.</p> <p>(18) The information system fails securely in the event of an operational failure of a boundary protection device.</p>
SC-26 <i>HONEYPOTS</i>	<p>Control: The information system includes components specifically designed to be the target of malicious attacks for the purpose of detecting, deflecting, and analyzing such attacks.</p> <p>Enhancement/s:</p> <p>(1) The information system includes components that proactively seek to identify web-based malicious code.</p>
SI-4 <i>INFORMATION SYSTEM MONITORING</i>	<p>Enhancement/s:</p> <p>(10) The organization makes provisions so that encrypted traffic is visible to information system monitoring tools.</p>

9 Directives, Policies, and Standards

This section identifies existing federal laws, Executive Orders, regulations, directives, policies, and standards applicable to the Capability but does not include those that are agency specific.

Network Boundary Protection Directives and Policies

Title, Date, Status	Excerpt / Summary
Intelligence Community (IC)	
Nothing found	
Comprehensive National Cybersecurity Initiative (CNCI)	
NSPD-54/HSPD-23 Cybersecurity Presidential Directive (Comprehensive	Summary: National Security Presidential Directive-54/Homeland Security Presidential Directive-23 (NSPD-54/HSPD-23), in which the Comprehensive National



CGS Network Boundary Protection Capability



Version 1.1.1

National Cybersecurity Initiative [CNCI]), 8 January 2008, Classified	Cybersecurity Initiative (CNCI) is described, is classified. Initiative 7 deals with increasing the security of classified networks.
Department of Defense (DoD)	
CJCSI 6510.01E, Information Assurance (IA) and Computer Network Defense, 12 August 2008, Unclassified	Summary: This instruction provides joint policy and guidance for information assurance (IA) and Computer Network Defense (CND) operations. Policy includes the following: interconnection of information systems will be managed to continuously minimize community risk and ensure that the protection of one system is not undermined by vulnerabilities of other interconnected systems. Firewalls, Cross-Domain Solutions, access control lists (ACLs), intrusion prevention systems, Demilitarized Zones (DMZs), and other protection procedures and devices will be used to restrict access to and from isolated local area network (LAN) segments.
DoD Firewall Guidance, version 1.2, 9 March 2001, Unclassified	Summary: This document addresses Department of Defense (DoD) firewall architecture deployment strategies and baseline firewall configurations for installing firewalls in classified or unclassified networks while supporting the Defense-in-Depth strategy. It includes firewall types, assurance levels, firewall services, and potential threats. Although not a policy document, it contains highly recommended guidance.
DISA Network Infrastructure Security Technical Implementation Guide (STIG), version 7.1, 25 October 2007, Unclassified	Summary: This Security Technical Implementation Guide (STIG) provides security considerations at the network level needed to achieve an acceptable level of risk for information as it is transmitted through an enclave. It was developed to enhance the confidentiality, integrity, and availability of sensitive DoD Automated Information Systems.
DISA Enclave Security Security Technical Implementation Guide (STIG), version 4.2, 10 March 2008, Unclassified	Summary: This STIG provides Organizations an overview of the applicable policy and additional STIG documents required to implement secure information systems and networks while ensuring interoperability.
DISA Internet-NIPRNet Department of Defense (DoD) Demilitarized Zone	Summary: Although this document is intended for use when implementing Internet/Nonclassified Internet Protocol Router Network (NIPRNet) DMZ connections, where appropriate,



CGS Network Boundary Protection Capability



Version 1.1.1

(DMZ) Engineering Plan, version 0.9, January 2010, Unclassified	the guidance can be applied to other security environments. This document provides an engineering plan to implement and manage DMZs.
Committee for National Security Systems (CNSS)	
Nothing found	
Other Federal (OMB, NIST, ...)	
Nothing found	
Executive Branch (EO, PD, NSD, HSPD, ...)	
Nothing found	
Legislative	
Nothing found	

Network Boundary Protection Standards

Title, Date, Status	Excerpt / Summary
Intelligence Community (IC)	
Nothing found	
Comprehensive National Cybersecurity Initiative (CNCI)	
Nothing found	
Department of Defense (DoD)	
Nothing found	
Committee for National Security Systems (CNSS)	
Nothing found	
Other Federal (OMB, NIST, ...)	
NIST SP 800-41, Guidelines on Firewalls and Firewall Policy, September 2009,	Summary: This special publication provides an overview of firewall technologies and discusses their security capabilities and relative advantages and disadvantages; it also provides examples of where firewalls can be placed within networks



CGS Network Boundary Protection Capability



Version 1.1.1

Unclassified	and the implications of deploying firewalls in particular locations; and recommendations for establishing firewall policies and for selecting, configuring, testing, deploying, and managing firewall solutions.
Executive Branch (EO, PD, NSD, HSPD, ...)	
Nothing found	
Legislative	
Nothing found	
Other Standards Bodies (ISO, ANSI, IEEE, ...)	
Nothing found	

10 Cost Considerations

This section provides examples of some of the types of costs that the Organization will need to consider when implementing this Capability. The following examples are costs that are common across all of the Community Gold Standards Capabilities:

1. Solution used for implementation (hardware and/or software)
2. Necessary training
3. Licensing (if applicable)
4. Lifecycle maintenance
5. Impact/dependency on existing services
6. Manpower to implement, maintain, and execute
7. Time to implement, maintain, and execute
8. Network bandwidth availability and consumption
9. Scalability of the solution relative to the Enterprise
10. Storage and processing requirements

In addition to the common costs, the following are examples of cost considerations that are specific to this Capability:

1. Not all data can be transferred because of different security levels—Some mission objectives may be hindered because the security classification of the data precludes sharing.



CGS Network Boundary Protection Capability



Version 1.1.1

2. Integration with existing, if any, protection systems—Layered defense is important, but if implemented incorrectly defenses can interfere with one another to reduce overall protection or unnecessarily constrict authorized traffic.
3. Extensibility/scalability of solution—As the network grows, the solution(s) will need to be able to scale to ensure it does not become a bottleneck in the future.
4. Necessary training—The Enterprise needs to consider whether the solution will be administered by another Organization (e.g., CDS Enterprise Service) or internally administered.

11 Guidance Statements

This section provides Guidance Statements, which have been extracted from Section 3 (Capability Gold Standard Guidance) of this Capability document. The Guidance Statements are intended to provide an Organization with a list of standalone statements that are representative of the narrative guidance provided in Section 3. Below are the Guidance Statements for the Network Boundary Protection Capability.

- The Enterprise shall protect and control access to Enterprise resources across a security boundary, which is when there is a separation of entities (systems, networks, enclaves, or Enterprises), which are governed by differing security policies or operate in a different threat environment. Protection of network boundaries is carried out by placing IA mechanisms between the internal system and the systems external to the security boundary. Examples of such IA mechanisms include, but are not limited to, CDSs, Controlled Interfaces, DMZs, VPNs, and encryption devices at the boundary, or interface.
- All network boundaries shall (based on connection type, mission needs, risk, etc.) incorporate approved firewalls, intrusion prevention/detection solutions, spam/malware/content filtering, DMZs for vulnerable border resources, and/or CDSs where the network connects to another network of a different security policy or environment.
- The Enterprise shall employ vendor diversity or defense-in-depth strategies for devices, where appropriate.
- The Enterprise shall verify all remote connections, individually, each time a connection attempt is made.
- Physical access to network resources shall also be secured (as deemed by the accrediting authority), including investigating users prior to granting access, authenticating users via secure credentials upon entry, employing guards, using security systems, etc.



CGS Network Boundary Protection Capability



Version 1.1.1

- All local resources shall be verified, local connections authorized, and WAPs secured.
- The boundary devices shall provide sufficient levels of protection at all layers (i.e., physical through application) to ensure the added risks are sufficiently mitigated.
- Processes shall be in place to monitor both the boundary defenses and all traffic flowing through the boundary.
- Network boundary devices shall detect events or incidents and forward the incident information to an incident response system.
- CDSs are placed at the boundary of enclaves with differing security domains; these devices shall be high-assurance devices (i.e., devices that have been carefully designed and implemented for high-risk situations) and undergo specialized certification testing.
- DMZs shall be placed as close to the network boundary as possible and implemented with specialized IA controls and devices that are specifically designed for the protocols or services used at the boundary connection.
- Remote access shall be provided by a secure service, such as a VPN, that provides two-way authentication between the remote user and local system and, at a minimum, encrypts the data when outside the network boundary.
- Encrypted traffic shall be passed only to allowed access points and the appropriate policies shall be defined for the encrypted traffic (e.g., it is dropped at the boundary, etc.). Network boundary protections shall differentiate between the type of access and functionality that shall be allowed with each remote connection.
- All boundary devices shall be capable of performing bidirectional information services or data flows.
- All boundary devices shall be capable of packet inspection and filtering at the IP (routes) and application (deep packet) layer.
- All boundary devices shall be capable of scanning for malicious content.
- All boundary devices shall be capable of authenticating traffic sources and destinations and enforcing the access policy.
- All boundary devices shall be capable of auditing, monitoring, and responding or providing information to the appropriate system to respond.
- All boundary devices shall be capable of meeting the Enterprise's throughput and availability requirements, providing mechanisms such as redundancy and failover, where appropriate.
- All boundary devices shall be capable of enforcing policies set by the Enterprise.
- There shall be trustworthy mechanisms to administer, manage, and monitor network boundary protection solutions.



CGS Network Boundary Protection Capability



Version 1.1.1

- All devices shall be on the evaluated and approved products list for the Organization and the specific environment where they are being implemented.